FORECASTING FINANCIAL TIME SERIES WITH SPECTRUM ANALYSIS

[1]PIERRE ROSTAN, [2]ALEXANDRA ROSTAN

[1,2]American University in Cairo, Department of Management, AUC Avenue, P.O. Box 74, New Cairo 11835, Egypt
E-mail: [1]rostan.pierre@gmail.com, [2]millelys@gmail.com

Abstract - In Physics, spectrum analysis models electrical, audio or seismic signals; in this paper, it is applied to financial time series forecasting. Any time series of financial assets may be decomposed in simpler signals called approximations and details in the framework of the one-dimensional discrete wavelet analysis. The simplified signals are recomposed after extension. The final output is the forecasted time series that are compared to observed data. Results show the pertinence of adding spectrum analysis to the battery of tools used by econometricians and quantitative analysts for the forecast of economic or financial time series.

Keywords - Derivatives pricing; spectrum analysis, wavelet analysis.

I. INTRODUCTION

The one-dimensional discrete wavelet analysis is applied to financial times series forecasting. The intuition behind the use of wavelet analysis applied to times series forecasting is simple. Many physical phenomena such as electrical, audio or seismic signals propagate through space in waveforms. In finance and economics, interest rates, exchange rates, volatility of asset returns, gross domestic product, levels of employment or consumer spending propagate through time in waveforms. For example, Rostan et al. (2016) have shown that interest rates may be forecasted using spectrum analysis for their valuable property of propagating through time in waveforms. Population estimates follow also waveform patterns and may be forecasted using spectrum analysis (Rostan et al., 2015). The challenge in spectrum analysis is to apply this analysis to asset prices time series when they do not follow a waveform pattern such as non-stationary S&P 500 index prices time series. The intuitive idea will be to decompose the original signal of S&P 500 index prices (i.e. the initial time series) in waveform signals, then to apply spectrum analysis to these decomposed signals since waveform patterns are easily captured by spectrum analysis.

Section 2 will present the methodology. Section 3 will present the results and section 4 discussion and conclusion.

II. METHODOLOGY

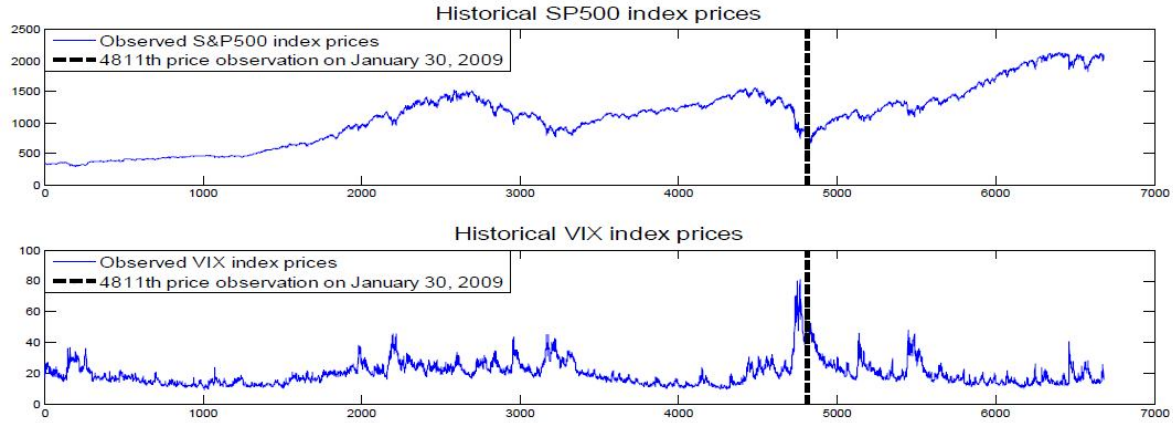
The objective of the paper is to apply spectrum analysis to the forecast of financial times series and to illustrate the versatility of the method with different types of times series.

We focus on two indices, one promoted, the other built by the Chicago Board of Options Exchange (CBOE) since CBOE lists index options. The two indices are the S&P 500 Composite index and the VIX index illustrated by Figure 1. The S&P 500 Composite index is a stock market index based on the market capitalizations of five hundred large cap companies listed on U.S. financial markets. This index is frequently used as benchmark of the U.S. stock market. The VIX index is built by CBOE from the implied volatility of S&P 500 index options and it gauges the market participants' expectation of the stock market volatility over the next 30 days.

The reason of choosing these two indices lies on the divergent properties of their time series: applying the Augmented Dickey-Fuller test on index prices of these two series, either on the full sample extending from January 2, 1990 to June 30, 2016 (6,678 data) or on the sub-sample extending from January 2, 1990 to June 30, 2009 (4,811 data) lead to the same results: S&P 500 Composite index series is non-stationary (i.e. existence of a unit root) and VIX index series is stationary (rejection of the existence of a unit root).
Forecasting with spectrum analysis non-stationary and stationary time series with success will illustrate the versatility of the spectrum analysis method for distinctive financial series.

To generate index forecasts, we need: 1) to decompose the initial signal, i.e the initial index time series, into decomposed signals; 2) to forecast the decomposed signals and to reconstruct the decomposed signals. These two steps are explained below.

Figure 1: Time series of historical S&P 500 Composite index and VIX index prices from January 2, 1990 to June 30, 2016 (6,678 data). Source: CBOE website, http://www.cboe.com/micro/vix/historical.aspx

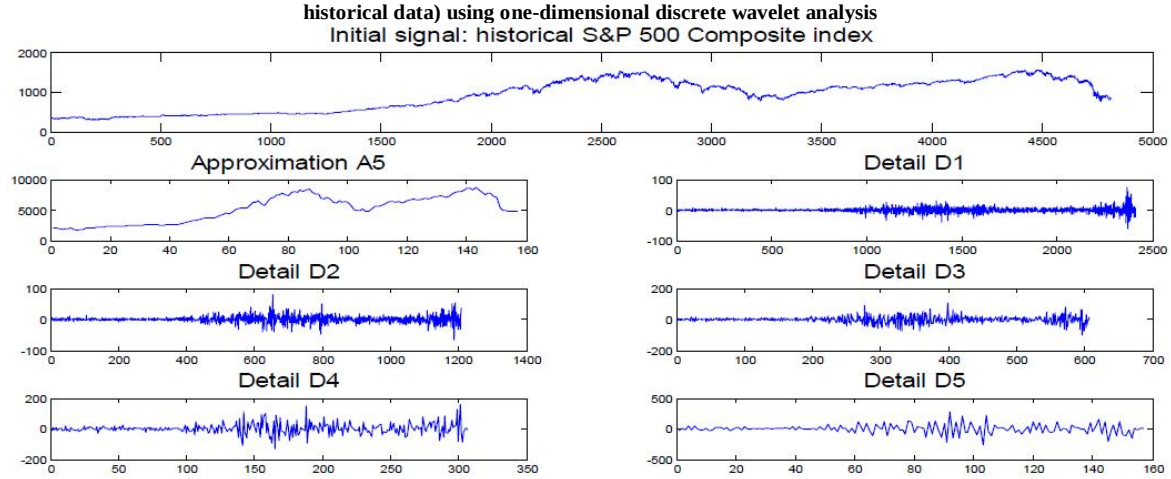


Step 1

Signal processing models wave-type signals studied in Physics such as electrical signals, audio signal processing or seismology. The VIX index (bottom graph of Figure 1) looks like a signal. The S&P 500 Composite index (top graph of Figure 1) does not clearly follow a waveform pattern. However, a sagacious observer may identify two cycles for the S&P 500 Composite index of about 6-year frequency between May 1997 (the 1,855th observation being May 1, 1997 on Figure 1) and January 2009 (the 4,811th observation being January 30, 2009), but thereafter the following uptrend departs from a waveform and behaves like a submartingale, i.e. a fair-game where S&P 500 index prices in the next period are expected to be greater than S&P 500 index prices in the current period. This is this last period of the S&P 500 Composite and VIX indices in Figure 1 starting at the vertical dot line that we try to capture and forecast with spectrum analysis.

The one-dimensional discrete wavelet analysis is the preferred method to decompose the initial signal. As an example, Figure 2 illustrates five levels of decomposition of the historical S&P 500 Composite index signal ranging from January 2, 1990 to June 30, 2009 (4,811 historical data).

Figure 2: Five-level decomposed signals of the S&P 500 Composite index prices from January 2, 1990 to June 30, 2009 (4,811 historical data) using one-dimensional discrete wavelet analysis



The one-dimensional initial signal S&P 500 Composite index is decomposed in simpler signals called approximations A and details D, such as: Initial Signal = A5 + D1 + D2 + D3 + D4 + D5.

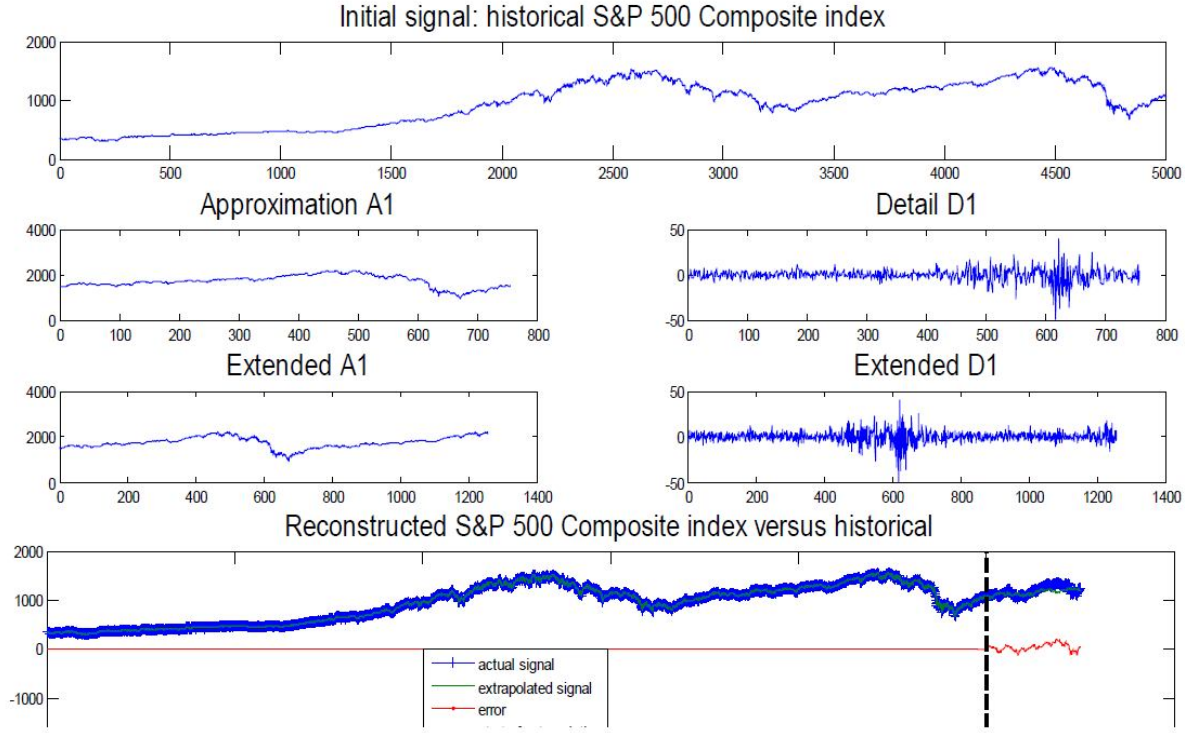
Step 2

For simplicity, we implement step 2 with the one-level decomposition/reconstruction method leading to one-level decomposed signals A1 and D1. After decomposition presented in step 1, we apply in step 2:

1) Wavelet extension to A1 and D1, then we reconstruct the signal from these two extended signals. Figure 3 illustrates the one-level decomposition of the initial signal S&P 500 Composite index in A1 and D1 from January 2, 1990 to October 29, 2009 (5,000 historical data), then its reconstruction after A1 and D1 have been extended with wavelet extension to forecast 500 data from October 30, 2009 to October 24, 2011. As observed, the extrapolated signal which is the forecasted S&P 500 Composite index matches pretty well the observed S&P 500 Composite index as illustrated by a low error of forecasting.

2) Burg extension to A1 and D1, then we reconstruct the signal from the extended two signals.

Figure 3: One-level decomposed signals A1 and D1 of the S&P 500 Composite index prices from January 2, 1990 to October 29, 2009 (5,000 historical data) using one-dimensional discrete wavelet analysis + Wavelet extension of A1 and D1 + Reconstruction of the extended signal from A1 and D1 and forecasting of 500 data. Size of historical data used in wavelet extension: 1,500.



To evaluate the forecasting ability of spectrum analysis for the two series, we will test 20-day, 60-day, 126-day, 252-day and 500-day forecasting periods in business days that roughly represent 1-month, 3-month, 6-month, 1-year and 2-year forecasting periods in calendar days. The forecasting periods are non-overlapping. We benchmark spectrum analysis to the ARIMA(1,1,1) forecasting model. We use the Root Mean Error Square Criteria (forecasts versus historical data) to compute the error of forecasting.

Our forecasting period starts on January 30, 2009 and extends to June 30, 2016 (1,867 business days).

III. RESULTS

For spectrum analysis, we first identify the optimal size of historical data that will be used to forecast a given period for wavelet extension and the optimal autoregressive pth order for Burg extension. For example, to identify the optimal historical size (autoregressive pth order) used in the 20-day forecast, we forecast the last 20 days in-sample with spectrum analysis using wavelet (Burg) extension from January 2, 2009 to January 30, 2009. We make the historical period (autoregressive pth order) before January 2, 2009 varying from 50 to 2,500 historical data (autoregressive pth order) by step of 50. We apply this methodology for each forecast horizon (20-, 60-, 126-, 252-, 500-day) and gather the optimal size of historical data (optimal autoregressive pth order) in table 1 based on the criteria of Sum of Squared Errors (forecasts versus historical data).

Table 1: Optimal size of historical data for each forecast horizon with spectrum analysis

Forecast horizon in business days:	20	60	126	252	500
Forecast horizon in calendar month/year:	1-month	3-month	6-month	1-year	2-year
S&P 500 Composite index					
Optimal size of historical data with Wavelet extension:	1,450	1,450	1,600	1,800	1,800
Optimal autoregressive p^{th} order with Burg extension:	50	2,500	2,500	1,700	100
VIX index					
Optimal size of historical data with Wavelet extension:	1,600	1,550	1,550	1,550	1,550
Optimal autoregressive p^{th} order with Burg extension:	2,050	1,550	1,550	1,400	950

Table 2 summarize the results for S&P 500 Composite index and VIX index time series forecasted with spectrum analysis and ARIMA(1,1,1).

Table 2: Average Root Mean Square Error (forecasts versus historical data) for forecasts with spectrum analysis and ARIMA(1,1,1) model applied to S&P 500 Composite index and VIX index

Forecast horizon in business days:	20	60	126	252	500	
Forecast horizon in calendar month/year:	1-month	3-month	6-month	1-year	2-year	
Number of non-overlapping forecasting periods over 1,867 days:	93	31	14	7	3	
Average RMSE for S&P 500 Composite index (forecast versus historical data):						Overall Average:
•Using spectrum analysis with Wavelet extension:	281	371	436	346	341	355
•Using spectrum analysis with Burg extension:	33	75	253	171	243	155
•Using ARIMA(1,1,1) model:	34	64	235	143	270	149
Average RMSE for VIX index (forecast versus historical data):						
•Using spectrum analysis with Wavelet extension:	8	11	12	14	16	12
•Using spectrum analysis with Burg extension:	3	6	7	7	8	6
•Using ARIMA(1,1,1) model:	3	4	6	7	10	6

DISCUSSION AND CONCLUSION

- Based on the overall average RMSE (last column of Table 2), ARIMA(1,1,1) is slightly better than spectrum analysis with Burg extension (149 versus 155) to forecast the S&P 500 Composite index and as good as spectrum analysis with Burg extension (6 versus 6) to forecast the VIX index.
- For both time series, spectrum analysis with wavelet extension underperform the two other models.
- Spectrum analysis with Burg extension is as good or better to forecast S&P 500 Composite index and VIX index than ARIMA(1,1,1) for the 20-day and 50-day forecast horizons.
- To illustrate the forecasting ability of ARIMA(1,1,1) and spectrum analysis, we graph a 20-day forecast on July 15, 2013. Looking at Figure 4, we undestand better why spectrum analysis with wavelet extension underperforms the two other models: its extension is more wacky than the other two. It is as if its waveform nature pressures the signal extension to look like a wave when it is not necessary.

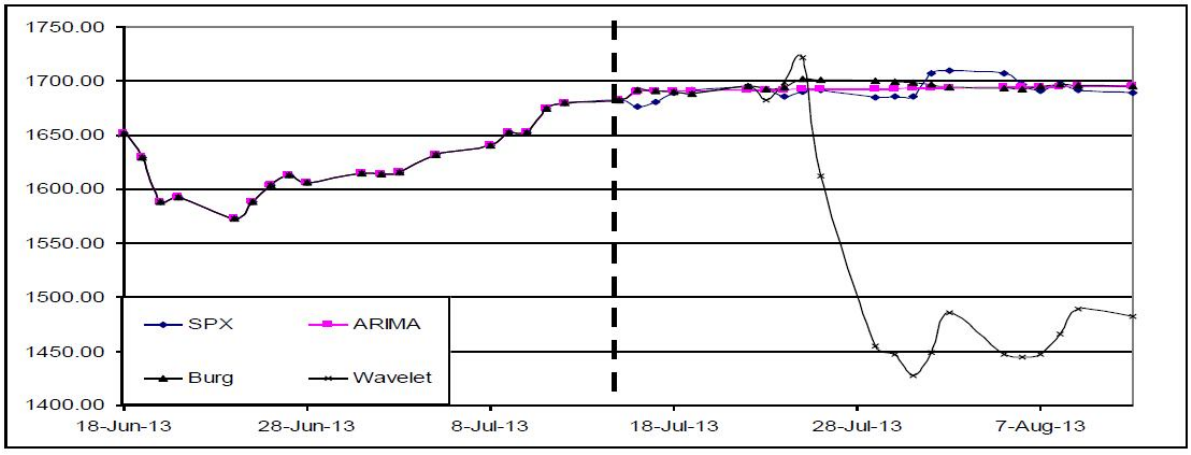


Figure 4: 20-day forecast of the S&P 500 Composite index on July 15, 2013 using ARIMA(1,1,1) and spectrum analysis with Wavelet and Burg extensions (the dot line represents the start of extrapolation)

In conclusion, we introduce spectrum analysis for the forecasting of financial time series. We measure its ability on two series with divergent properties, one series being stationary (the VIX index), the other one being non-stationary (the S&P 500 Composite index). We show that overall for 5 different forecast horizons (20-, 60-, 126-, 252-, 500-day), spectrum analysis with Burg extension follows closely the ARIMA(1,1,1) model for its ability to forecast the S&P 500 Composite index and is as good as ARIMA(1,1,1) to forecast the VIX index. Thus, we illustrate the versatility of the method to forecast distinctive times series. We also demonstrate that the nature of the extension of the decomposed signals is important since wavelet extension underperforms Burg extension. In this paper, we have tested one-level decomposition of signals. Looking at multi-

level decompositions or further extensions open ways to new researches.

REFERENCES

[1] Baillie, R., and T. Bollerslev. "Prediction in Dynamic Models with Time-Dependent Conditional Variances." Journal of Econometrics. Vol. 52, 1992, pp. 91–113.

[2] Bollerslev, T. "Generalized Autoregressive Conditional Heteroskedasticity." Journal of Econometrics. Vol. 31, 1996, pp. 307–327.

[3] Bollerslev, T. "A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return." The Review Economics and Statistics. Vol. 69, 1987, pp. 542–547.

[4] Box, G. E. P. and Jenkins, G. M. (1976): "Time Series Analysis: Forecasting and Control", Revised Edition, San Francisco: Holden Day.

[5] Box, G. E. P., G. M. Jenkins, and G. C. Reinsel. Time Series Analysis: Forecasting and Control 3rd ed. Englewood Cliffs, NJ: Prentice Hall, 1994.

[6] Burg, J.P. (1975): "Maximum Entropy Spectral Analysis". Available at:

[7] http://sepwww.stanford.edu/theses/sep06/ [September 10, 2016]

[8] Durbin, J. (1960): "The fitting of time series models", Revue de l'Institut International de Statistique, 28, pp. 233-44.

[9] Enders, W. Applied Econometric Time Series. Hoboken, NJ: John Wiley & Sons, 1995.

[10] Engle, R. F. "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation." Econometrica. Vol. 50, 1982, pp. 987–1007.

[11] Hamilton, J. D. "Time Series Analysis." Princeton, NJ: Princeton University Press, 1994.

[12] Kay, S.M. "Modern Spectral Estimation: Theory and Application". Englewood Cliffs, NJ: Prentice Hall, 1988, pp. 228–230.

[13] Levinson, N. (1946): "The Wiener RMS (root mean square) error criterion in filter design and prediction", Journal of Mathematical Physics, 25, pp. 261-78.

[14] Rostan, P., Belhachemi, R., & Rostan, A. (2015). "Appraising the financial sustainability of a pension system with signal processing." Estudios De Economía Aplicada, 33, 1-16.

[15] Rostan, P., Belhachemi, R., & Racicot, F.E. (2016). "Forecasting the yield curve with the Burg model." Journal of Forecasting, doi: 10.1002/for.2416

[16] Tuljapurkar, S. and Boe, C (1999): "Validation, probability-weighted priors, and information in stochastic forecasts", International Journal of Forecasting, 15, pp. 259–271.

★★★